

# RankMerging: A supervised learning-to-rank framework to predict links in large social networks

Lionel Tabourier<sup>\*1</sup>, Daniel F. Bernardes<sup>2</sup>, Anne-Sophie Libert<sup>3</sup>,  
and Renaud Lambiotte<sup>3</sup>

<sup>1</sup>CNRS, UMR 7606, LIP6,, Sorbonne Universités, UPMC  
Université Paris 06, UMR 7606, LIP6, 75252 Paris, France

<sup>2</sup>L2TI, Université Paris-Nord, France

<sup>3</sup>naXys, University of Namur, Namur, Belgium

## Abstract

Uncovering unknown or missing links in social networks is a difficult task because of their sparsity and because links may represent different types of relationships, characterized by different structural patterns. In this paper, we define a simple yet efficient supervised learning-to-rank framework, called *RankMerging*, which aims at combining information provided by various unsupervised rankings. We illustrate our method on three different kinds of social networks and show that it substantially improves the performances of unsupervised metrics of ranking. We also compare it to other combination strategies based on standard methods. Finally, we explore various aspects of *RankMerging*, such as feature selection and parameter estimation and discuss its area of relevance: the prediction of an adjustable number of links on large networks.

---

<sup>\*</sup>corresponding author: [lionel.tabourier@ens-lyon.org](mailto:lionel.tabourier@ens-lyon.org)

# 1 Introduction

Link prediction is a key field of research for the mining and analysis of large-scale social networks because of its many practical applications: going from recommendation strategies for commercial websites [16] to recovering missing links in incomplete data [34]. Link prediction also has significant implications from a fundamental point of view, as it allows for the identification of the elementary mechanisms behind the creation and decay of links in time-evolving networks [19]. For example, triadic closure, at the core of standard methods of link prediction is considered as one of the driving forces for the creation of links in social networks [18].

In their seminal formulation, Liben-Nowell and Kleinberg [20] present link prediction as follows: considering a snapshot of a network at time  $t$ , the problem is to predict which links will be present at a future time  $t' > t$ . It is therefore a binary classification issue, where features used are mostly based on the structural properties of the network or nodes and links attributes. Our problem is similar, except for the fact that there is no timescale involved, we aim at recovering missing links in large social networks. In this context, the classification issue has specific characteristics: for a typical  $10^6$  nodes network, there are around  $10^{12}$  candidate pairs of nodes that can be connected, most of them being completely irrelevant. The problem is unmanageable without restraining ourselves to subsets of pairs, but even with this restriction we have to handle typically  $10^5$  to  $10^7$  items, as we shall see it implies solving challenges in terms of computational efficiency. Features are indeed known to be domain-specific: as links play different roles in social networks, they are expected to be surrounded by different types of environments and thus to be best identified by different features. For these reasons, schemes based on a single metric are prone to misclassification. In the context of classification, machine learning methods have been widely used to combine the available information. In recent works, classification trees, support vector machines, matrix factorization or neural networks are implemented to predict links in biological networks or scientific collaboration networks [27, 17, 3, 21, 24, 11]. However, these classification methods are not designed to easily set the number of predictions, while this property is desirable in the context of link prediction in social networks.

Another way to address the issue consists in establishing a ranking of likely links according to a scalar metric, correlated with the existence of interactions between nodes. The user sets the number of links predicted by

selecting the  $T$  top-ranked pairs. Ranking metrics are often based on the structural properties of the network of known interactions, either at a local scale – e.g. the number of common neighbors – or at a global scale – e.g. random walk or hitting time. See for example [23] for a survey. Other sources of information are available to predict links, in particular node attributes such as age, gender or other profile information [2, 4], geographic location [31], as well as interaction attributes: frequencies [33] or the time elapsed since the last interaction [30].

In this context, *learning-to-rank* frameworks can be used to combine different metrics for link prediction. Unsupervised solutions are available, such as Borda’s method or Markov chain ordering [13, 32]. But formulating the link prediction problem in a supervised way should perform better, so we explore supervised methods available for learning-to-rank purpose. Most supervised learning-to-rank techniques have been designed in the context of information retrieval tasks, such as document filtering, spam webpage detection, recommendation or text summarization, e.g. [14, 5, 9]. In [22], the author distinguishes between three kinds of approaches in this field: pointwise approaches are the most straightforward, using the score or rank associated to a feature and using it to fit, for example, a regression model. One undesirable effect is that low-ranked items tend to have an over-important role in the learning process, which is particularly critical in the case of link prediction as rankings are very large. Pairwise approaches [15] consist in transforming the ranking problem into a classification one, by considering couple of items and learning which one should be ranked above. This transform allows the use of many supervised classification methods, but unfortunately even the cheapest implementations of this approach [8] cannot be used to predict links on large networks, as the number of items to rank here is larger than  $10^6$ . Listwise approaches [6] use a ranking of items as ground truth, which is not relevant to our case: when predicting  $T$  links, the quality of the prediction provided by two rankings is strictly equivalent if they provide the same amount of true prediction in their top- $T$  items. More generally, information retrieval techniques primarily aim at high precision on the top-ranked items, and stress the relative ranking of two items. As stated in [7], most of the research on the topic has therefore focused on improving the prediction accuracy rather than making the algorithms scalable, which is crucial in the case of link prediction. Learning-to-rank in the context of link prediction in large graphs calls for specific methods, suited for large rankings. Closer to our work, in [29] links are predicted in scientific collaboration networks using

supervised approaches based on improvements of unsupervised methods, it allowed efficient predictions on  $10^4$  to  $10^5$  items rankings.

Here, we propose a simple yet efficient learning-to-rank supervised framework specifically designed to uncover links in large and sparse networks, such as social networks. We improve the prediction by combining rankings obtained from different sources of information. The article is organized as follows. In Section 2, we describe the datasets under study in this article, which are three large social networks (featuring more than  $10^6$  nodes) representing different kinds of social interactions: phone calls, coauthorship, and friendship on an online social network. Section 3 is dedicated to the description of the metrics that we use to evaluate the performances of the link prediction. We then present in Section 4 how classic unsupervised learning methods can be applied to the problem under consideration. In Section 5, we develop a supervised machine learning framework, called *RankMerging*, which improves the quality of predictions by aggregating the information from the unsupervised metrics. Finally, we implement this method on the three datasets during two series of experiments in Section 6, comparing our results to those of other methods. We also explore aspects such as the feature selection problem, the impact of parameters values etc, and show that *RankMerging* is suited to social networks where information is partial and noisy, and the number of links to predict is large.

## 2 Datasets

### 2.1 PSP phonecall network

We investigate a call detail record of approximately  $14 \cdot 10^6$  phonecalls of anonymized subscribers of a European phone service provider (PSP) during a one month period. Users are described as nodes and an interaction between two users is a directed link. We focus on the social groups underlying the phone call network, so in order to filter out calls which are not indicative of a lasting social relationship, we only consider calls on bidirectional links, i.e. links that have been activated in both directions. After this filtering has been applied, interactions between users are considered as undirected. The total number of phone calls between nodes  $i$  and  $j$  is the weight of this link. The resulting network is composed of 1,131,049 nodes, 795,865 links and 10,934,277 calls.

## 2.2 DBLP coauthorship network

The second dataset under study is the collection of computer science publications DBLP<sup>1</sup>, which contains roughly  $19 \cdot 10^6$  papers. We study the collaboration network, where nodes of the network are authors and links connect authors who are coauthors, moreover links are weighted by the number of collaborations. The network features 10,724,828 links and 1,314,050 nodes.

## 2.3 Pokec social network

Pokec is a popular online social network in Slovakia<sup>2</sup>. We consider the network of friendship links between accounts. In its original version there are around  $31 \cdot 10^6$  friendship links, but friendship is directed in this dataset, so we only kept reciprocated friendships in the spirit of the preprocessing of the PSP dataset. The final network contains 8,320,600 links and 1,632,804 nodes.

## 3 Performance evaluation

Link prediction is here formulated as a learning-to-rank problem. The definition of a *good* prediction and thus of an adequate quality estimator depends on the purpose of the ranking problem. For example, the goal of an efficient search engine is to provide highly relevant information on a small amount of items. More generally, in the field of information retrieval, the performance of a ranking is often measured using metrics that emphasize the relative ranking between two items, such as discounted cumulative gain. However, considering link prediction for a fixed number of predictions  $T$ , the link is predicted or not depending on whether its rank falls above or below  $T$ . The quality of a prediction is therefore assessed in this work by measuring the numbers of true and false positive (resp.  $\#tp$  and  $\#fp$ ), true and false negative ( $\#tn$  and  $\#fn$ ) predictions in the top  $T$  pairs, and usual related quantities: precision  $\mathbf{Pr} = \frac{\#tp}{\#tp + \#fp}$ , recall  $\mathbf{Rc} = \frac{\#tp}{\#tp + \#fn}$  and F-score  $\mathbf{F} = \frac{2 \cdot \mathbf{Pr} \cdot \mathbf{Rc}}{\mathbf{Pr} + \mathbf{Rc}}$ .

Previous works have emphasized the dramatic effect of class imbalance (or skewness) on link prediction problems in social networks, especially in mobile phone networks [21, 9]. The fact that the network is sparse and that there are

---

<sup>1</sup>Available on [http://konect.uni-koblenz.de/networks/dblp\\_coauthor](http://konect.uni-koblenz.de/networks/dblp_coauthor).

<sup>2</sup>Available on <http://snap.stanford.edu/data/soc-pokec.html>.

many more pairs of nodes than links makes the prediction and its evaluation tricky: the typical order of magnitude of the classes ratio for a social network made of  $N$  nodes is  $1/N$ . It means that the number of predicted links is much lower than the number of candidate pairs, consequently the fall-out  $\frac{\#fp}{\#fp+\#tn}$  is always very small, making the ROC curve an inappropriate way of visualizing the performances in our case. For this reason and because we aim at improving both precision and recall over a large range, we visualize the performances in the precision-recall space.

## 4 Unsupervised rankings

### 4.1 Ranking metrics

In this work, we focus on structural features that assign to each pair of nodes a score based on topological information, then pairs are ranked according to this score. A large number of metrics have been used in the past, see for example [34, 23]. The goal of this paper is not to propose elaborate classifiers, but to present a method that takes advantage of how complementary they are. We have therefore chosen classic metrics and generalized them to the case of weighted networks – other generalizations exist in the literature, e.g. [25].

#### 4.1.1 Local features

In the following,  $\mathcal{N}(i)$  denotes the set of neighbors of node  $i$ , its degree is  $\delta(i) = |\mathcal{N}(i)|$ ,  $w(i, j)$  is the weight of a link  $(i, j)$  and  $W(i)$  is the activity of a node  $i$ , that is the sum of the weights of its adjacent links. Some metrics are local (also called *neighborhood rankers*) as they only rank links among nodes which are at most at distance 2.

- *Common Neighbors index (CN)*, based on the number of common neighbors shared by nodes  $i$  and  $j$ , the corresponding unweighted and weighted scores are:

$$s_{CN}(i, j) = |\mathcal{N}(i) \cap \mathcal{N}(j)| \quad \text{and} \quad s_{CN_w}(i, j) = \sum_{k \in \mathcal{N}(i) \cap \mathcal{N}(j)} w(i, k) \cdot w(j, k)$$

- *Resource Allocation index (RA)*:

$$s_{RA}(i, j) = \sum_{k \in \mathcal{N}(i) \cap \mathcal{N}(j)} \frac{1}{\delta(k)} \quad \text{and} \quad s_{RA_w}(i, j) = \sum_{k \in \mathcal{N}(i) \cap \mathcal{N}(j)} \frac{1}{W(k)}$$

- *Adamic-Adar index (AA)*:

$$s_{AA}(i, j) = \sum_{k \in \mathcal{N}(i) \cap \mathcal{N}(j)} \frac{1}{\log(\delta(k))} \quad \text{and} \quad s_{AA_w}(i, j) = \sum_{k \in \mathcal{N}(i) \cap \mathcal{N}(j)} \frac{1}{\log(W(k))}$$

- *Sørensen index (SR)*:

$$s_{SR}(i, j) = \frac{2 \cdot |\mathcal{N}(i) \cap \mathcal{N}(j)|}{\delta(i) + \delta(j)} \quad \text{and} \quad s_{SR_w}(i, j) = \frac{\sum_{k \in \mathcal{N}(i) \cap \mathcal{N}(j)} w(i, k) + w(j, k)}{W(i) + W(j)}$$

#### 4.1.2 Global features

Another class of features is global, since they are calculated using the whole structure of the network, and allow for the ranking of distant pairs of nodes:

- *Katz index (Katz)*, computed from the number of paths from node  $i$  to node  $j$  of length  $l$ , i.e.  $\nu_{ij}(l)$ , according to the following expression ( $\beta$  is a parameter lower than 1):

$$s_{Katz}(i, j) = \sum_{l=1}^{\infty} \beta^l \nu_{ij}(l)$$

Note that in the weighted case, the number of paths is computed as if links were multilinks.

- *Random Walk with Restart index (RWR)*, derived from the *PageRank* algorithm,  $s_{RWR_w}(i, j)$  is defined as the probability that a random walker starting on node  $i$ , going from a node  $k$  to a node  $k'$  with probability  $p \cdot w(k, k')/W(k)$  and returning to  $i$  with probability  $1 - p$ , is on  $j$  in the steady state of the process.
- *Preferential Attachment index (PA)*, based on the observation that active nodes tend to connect preferentially in social networks.

$$s_{PA_w}(i, j) = W(i) \cdot W(j)$$

### 4.1.3 Intermediary features

In practice, the exact computation of global metrics is expensive on large networks, that is why approximations are often favoured to compute these scores. Both *Katz* and *RWR* are computed using infinite sums, in the following, it will be approximated by keeping only the four first dominating terms to reduce the computational cost. This approximation means that we can only predict links between pairs of nodes at a maximum distance of 4. Notice that it is a way to reduce the class-imbalance problem: as distant pairs are less likely to be connected, we dismiss them in order to increase the (true positive / candidate pairs) ratio. As the class imbalance problem is known to hinder dramatically the performance of PA, we have restricted the ranking in this case to pairs of nodes at a maximum distance of 3. Notice that with larger maximum distances, we can increase the maximum recall that can be reached, but at the cost of a drop of precision.

When even distance 4 approximation is too expensive, we use the *Local Path* index, especially designed to capture the structure at an intermediary scale:

$$s_{LP}(i, j) = \nu_{ij}(2) + \gamma \cdot \nu_{ij}(3)$$

With the same notations as *Katz* index and  $\gamma$  is a parameter.

## 4.2 Borda's method

The main purpose of this work is to develop a framework to exploit a set of  $\alpha$  rankings for link prediction. Here, we present an unsupervised way of merging rankings stemming from social choice theory: Borda's method is a *rank-then-combine* method originally proposed to obtain a consensus from a voting system [12]. Each pair is given a score corresponding to the sum of the number of pairs ranked below, that is to say:

$$s_B(i, j) = \sum_{\kappa=1}^{\alpha} |r_{\kappa}| - r_{\kappa}(i, j)$$

where  $|r_{\kappa}|$  denotes the number of elements ranked in  $r_{\kappa}$ . This scoring system may be biased toward global predictors by the fact that local rankings feature less elements. To alleviate this problem, pairs not ranked in ranking  $r_{\kappa}$  but ranked in  $r_{\kappa'}$  will be considered as ranked in  $r_{\kappa}$  on an equal footing



as any other unranked pair and below all ranked pairs. Borda’s method is computationally cheap, which is a highly desirable property in the case of large networks. A more comprehensive discussion on this method can be found in [13].

## 5 *RankMerging* framework

The ranking methods presented in the previous section use structural information in complementary ways. In systems such as social networks, where communication patterns of different groups, e.g. family, friends, coworkers etc, are different, one expects that a link detected as likely by using a specific ranking method may not be discovered using another one. In this section, we describe the supervised machine learning framework that we created to aggregate information from various ranking techniques for link prediction in social networks. In a nutshell, it does not demand for a pair to be highly ranked according to all criteria (as in a consensus rule), but at least one. The whole procedure is referred to as *RankMerging*<sup>3</sup>.

### 5.1 The method

#### 5.1.1 Learning phase.

We first consider the training set to learn the parameters on the learning graph  $\mathcal{G}_{learn}$ . Suppose that we have  $\alpha$  different rankings  $r_1, \dots, r_\alpha$  from the unsupervised methods described previously. During the learning phase, ranked pairs of nodes are labelled, depending on the fact that the link does exist or not. We create during the learning phase an output ranking  $r_{mer}^L$  which contains pairs selected from various rankings  $r_i$ .

**Sliding index.** Let  $\phi_i$  denote the index used to go through ranking  $r_i$ , we name it *sliding index* of  $r_i$  in the following. Predicted links on  $\mathcal{G}_{learn}$  rank from position 1 to  $\phi_i$  of each ranking  $r_i$ . We now consider simultaneously the  $\alpha$  rankings, with the goal to compute the optimal values of  $\phi_1, \dots, \phi_\alpha$  for a fixed number of predictions  $T$  so that the total number of true positive predictions  $\mathcal{S}(\phi_1, \dots, \phi_\alpha)$  is maximum. These values are the essential outputs

---

<sup>3</sup>An implementation and user guide are available on <http://lioneltabourier.fr/program.html>.

of the training phase, as they indicate the contribution of each ranking to the merged ranking. It is important to note that a link can only be predicted once: if a pair has been added to  $r_{mer}^L$  using ranking  $r_i$ , it cannot be added to  $r_{mer}^L$  anymore. It implies that we have to take into account the links already added to  $r_{mer}^L$  in the next steps, which makes this problem hard to solve exactly. For this reason, it is solved heuristically, as described below.

**Rationale.** The central idea is to find at each step the ranking with the highest number of true predictions in the next  $g$  coming steps. For that purpose, we define the *window*  $\mathcal{W}_i$  as the set of links predicted according to ranking  $r_i$  in the next  $g$  steps. Note that links already in  $r_{mer}^L$  are not considered for  $\mathcal{W}_i$ . The number of *tp* in  $\mathcal{W}_i$  is the *quality* of the window, denoted  $\chi_i$ . The ranking corresponding to the highest quality value is selected<sup>4</sup>, and we add its next top-ranked pair to the output ranking of the learning phase  $r_{mer}^L$ . Throughout the process, the sliding indices  $\phi_1 \dots \phi_\alpha$  are registered. Then, the windows  $\mathcal{W}_i$  are updated so that each one contains exactly  $g$  pairs, and the process is iterated until  $r_{mer}^L$  contains the chosen number of links to predict  $T$ . To summarize, we are looking for a maximum number of true predictions  $\mathcal{S}$  by local search. The algorithm corresponding to this process is described in Algo. 1.

**Illustration.** Table 1 gives an example of the first first steps of the merging process between two rankings  $r_A$  and  $r_B$  with  $g = 5$ . Pairs in the windows are represented with a gray background. Initially, there are 4 *tp* in  $\mathcal{W}_A$  and 3 *tp* in  $\mathcal{W}_B$ , consequently  $\chi_A > \chi_B$ , so the first link selected is the top-ranked pair available in  $\mathcal{W}_A$ : (1, 2). This pair is therefore excluded from the ranking  $r_B$ . At step 2, we have  $\chi_A = \chi_B = 4$ , the ranking with highest quality is then selected randomly, we suppose here that  $r_B$  has been selected so that the next link added to  $r_{mer}^L$  is (5, 18). At step 3,  $\chi_A = 4$  and  $\chi_B = 3$  then the pair selected to join  $r_{mer}^L$  is (1, 4) from  $r_A$ . At step 4,  $\chi_A = 4$  and  $\chi_B = 3$  and the pair selected is (5, 6) from  $r_A$ . At this point, according to the previously defined notations,  $\phi_1(4) = 3$  and  $\phi_2(4) = 1$ .

---

<sup>4</sup>In the case of a tie, we choose randomly.

$r_A$	$tp$	$r_B$	$tp$		$r_A$	$tp$	$r_B$	$tp$		$r_A$	$tp$	$r_B$	$tp$		$r_A$	$tp$	$r_B$	$tp$	
(1,2)		(5,18)	x		(1,2)		(5,18)	x		(1,2)		(5,18)	x		(1,2)		(5,18)	x	
(1,4)	x	(1,2)			(1,4)	x	(1,2)			(1,4)	x	(1,2)			(1,4)	x	(1,2)		
(5,6)	x	(8,9)			(5,6)	x	(8,9)			(5,6)	x	(8,9)			(5,6)	x	(8,9)		
(6,12)	x	(5,6)	x		(6,12)	x	(5,6)	x		(6,12)	x	(5,6)	x		(6,12)	x	(5,6)	x	
(5,18)	x	(7,11)	x		(5,18)	x	(7,11)	x		(5,18)	x	(7,11)	x		(5,18)	x	(7,11)	x	
(3,4)		(6,9)	x		(3,4)		(6,9)	x		(3,4)		(6,9)	x		(3,4)		(6,9)	x	
(4,9)	x	(1,14)			(4,9)	x	(1,14)			(4,9)	x	(1,14)			(4,9)	x	(1,14)		
(7,11)	x	(2,9)			(7,11)	x	(2,9)			(7,11)	x	(2,9)			(7,11)	x	(2,9)		
(2,9)		(3,7)			(2,9)		(3,7)			(2,9)		(3,7)			(2,9)		(3,7)		

$r_A$	$tp$	$r_B$	$tp$		$r_A$	$tp$	$r_B$	$tp$		$r_A$	$tp$	$r_B$	$tp$		$r_A$	$tp$	$r_B$	$tp$	
(1,2)		(5,18)	x		(1,2)		(5,18)	x		(1,2)		(5,18)	x		(1,2)		(5,18)	x	
(1,4)	x	(1,2)			(1,4)	x	(1,2)			(1,4)	x	(1,2)			(1,4)	x	(1,2)		
(5,6)	x	(8,9)			(5,6)	x	(8,9)			(5,6)	x	(8,9)			(5,6)	x	(8,9)		
(6,12)	x	(5,6)	x		(6,12)	x	(5,6)	x		(6,12)	x	(5,6)	x		(6,12)	x	(5,6)	x	
(5,18)	x	(7,11)	x		(5,18)	x	(7,11)	x		(5,18)	x	(7,11)	x		(5,18)	x	(7,11)	x	
(3,4)		(6,9)	x		(3,4)		(6,9)	x		(3,4)		(6,9)	x		(3,4)		(6,9)	x	
(4,9)	x	(1,14)			(4,9)	x	(1,14)			(4,9)	x	(1,14)			(4,9)	x	(1,14)		
(7,11)	x	(2,9)			(7,11)	x	(2,9)			(7,11)	x	(2,9)			(7,11)	x	(2,9)		
(2,9)		(3,7)			(2,9)		(3,7)			(2,9)		(3,7)			(2,9)		(3,7)		

$r_A$	$tp$	$r_B$	$tp$		$r_A$	$tp$	$r_B$	$tp$	
(1,2)		(5,18)	x		(1,2)		(5,18)	x	
(1,4)	x	(1,2)			(1,4)	x	(1,2)		
(5,6)	x	(8,9)			(5,6)	x	(8,9)		
(6,12)	x	(5,6)	x		(6,12)	x	(5,6)	x	
(5,18)	x	(7,11)	x		(5,18)	x	(7,11)	x	
(3,4)		(6,9)	x		(3,4)		(6,9)	x	
(4,9)	x	(1,14)			(4,9)	x	(1,14)		
(7,11)	x	(2,9)			(7,11)	x	(2,9)		
(2,9)		(3,7)			(2,9)		(3,7)		

Table 1: Four steps illustrating the learning algorithm with rankings  $r_A$  and  $r_B$  ( $g = 5$ ). Pairs predicted (i.e.  $\in r_{mer}^L$ ) have green backgrounds. Pairs with gray backgrounds are in the windows  $\mathcal{W}_A$  and  $\mathcal{W}_B$ . Barred pairs with red backgrounds have already been predicted and cannot be selected anymore.

### 5.1.2 Test phase.

**Rationale.** The test phase of the procedure consists in combining rankings on the test network  $\mathcal{G}_{test}$  according to the  $\phi_1 \dots \phi_\alpha$  learnt on the training network  $\mathcal{G}_{learn}$ . The practical implementation is simple: at each step, we



described in Algo. 2. Note that it is possible that the rankings considered during the learning phase have different length from the ones during the test phase. There are ways to circumvent this problem, namely to use a scaling factor  $f = |r^T|/|r^L|$ , that is the ratio of sizes between these rankings, so that considering step  $n$  of the merging algorithm, we use the ranking predicted on step  $s = \lceil \frac{n}{f} \rceil$  of the learning algorithm.

learning step $s$		$\phi_A = \phi[1]$	$\phi_B = \phi[2]$
1		1	0
2		1	1
3		2	1
4		3	1

$r_A$	$r_B$		$r_A$	$r_B$		$r_A$	$r_B$		$r_A$	$r_B$
(2,8)	(1,8)		(2,8)	(1,8)		(2,8)	(1,8)		(2,8)	(1,8)
(1,8)	(9,11)	→	(1,8)	(9,11)	→	(1,8)	(9,11)	→	(1,8)	(9,11)
(5,11)	(4,5)		(5,11)	(4,5)		(5,11)	(4,5)		(5,11)	(4,5)
(3,6)	(5,11)		(3,6)	(5,11)		(3,6)	(5,11)		(3,6)	(5,11)

Table 2: Illustration of the test algorithm (with  $f = 1$ ). Top:  $\phi_i$  learnt during the learning phase. Bottom: pairs in  $r_{mer}^T$  at each step of the aggregation (green background), barred pairs with red backgrounds have already been predicted and cannot be selected anymore.

**Illustration.** On Table 2, we give an example of test using the  $\phi_i$  learnt on the example of Tab. 1 (with  $f = 1$ ). At the learning step 1,  $\phi_1(1) = 1$  and  $\phi_2(1) = 0$ , so the top-ranked item of  $r_A$  is selected at the testing step 1 to be in  $r_{mer}^T$ . We iterate this process: at step 2,  $\phi_1(2) = 1$  and  $\phi_2(2) = 1$ , we therefore add the top item of  $r_2$  to  $r_{mer}^T$ , etc.

### 5.1.3 Complexity

An important benefit of our learning algorithm is that we need to go through each learning ranking only once. Moreover, by using appropriate data structures to store the windows  $\mathcal{W}_i$  – e.g., associative arrays – it is easy to manage each update in  $O(1)$ . So if we have  $\alpha$  rankings and  $T$  predictions, it implies a  $O(\alpha T)$  temporal complexity. Similarly, the test phase demands to go through the test rankings once, yielding a  $O(\alpha T)$  complexity. Besides, most of the memory used is due to the storage of the rankings, which is also  $O(\alpha T)$ .

---

**Algorithm 2:** *RankMerging* method: testing algorithm.

---

**inputs** : table of rankings  $\mathcal{R}$  ( $r_i = \mathcal{R}[i]$ ); sliding index matrix  $\phi$ ;  
maximum number of predictions  $T$ ;  
**outputs**: merged ranking  $r_{mer}^T$ ;  
// initialization:  
**begin**  
     $\forall i \ c[i] \leftarrow 1;$                       // cursors of the sliding index matrix  
     $\forall i \ \tau[i] \leftarrow 0;$                     // number of actual predictions from  $r_i$   
     $n \leftarrow 1;$                                 // number of predictions  
**while**  $n \leq T$  **do**  
     $\forall i$  **if**  $\tau[i] < \phi[i][s]$  **then**  
        **if**  $\mathcal{R}[i][c[i]] \notin r_{mer}^T$  **then**  
             $r_{mer}^T[n] \leftarrow \mathcal{R}[i][c[i]];$   
             $\tau[i] \leftarrow \tau[i] + 1;$   
             $c[i] \leftarrow c[i] + 1;$   
         $n \leftarrow n + 1;$

---

These time and space consumptions are in general insignificant w.r.t. the complexities of the preliminary unsupervised classification methods.

## 5.2 Optimality justification

In this part, we demonstrate that the aggregation method that we used in *RankMerging* provides the maximum value for a function  $\mathcal{S}_q(\phi_1, \dots, \phi_\alpha)$ , close to the function  $S$  defined in Section 5.1.1, and we discuss numerically how close  $\mathcal{S}_q$  and  $S$  can be.

### 5.2.1 Sufficient condition to maximize $\mathcal{S}_q$

First, we consider a slightly different theoretical problem: suppose that we have  $\alpha$  functions  $q_1, \dots, q_\alpha$  with values from  $\mathbb{N}$  to  $[0; 1]$ . We aim at finding  $\phi_1, \dots, \phi_\alpha$  that maximize  $\mathcal{S}_q(\phi_1, \dots, \phi_\alpha) = \sum_{t_1=1}^{\phi_1} q_1(t_1) + \dots + \sum_{t_\alpha=1}^{\phi_\alpha} q_\alpha(t_\alpha)$ , with fixed  $T = \phi_1 + \dots + \phi_\alpha$ .

**Lemma.** *If  $\forall k, q_k$  is a decreasing function, then by adding  $\max_k \{q_k(t_k)\}$  at each step, we obtain the maximum value of  $\mathcal{S}_q$  for a fixed  $T$ .*

*Proof.* By recursion. For  $T = 1$ , the lemma is trivial. We then use a *reductio ad absurdum* argument: suppose  $\mathcal{S}_q(\phi_1, \dots, \phi_\alpha)$  is the maximum value for  $T$ , and let  $j$  be the index corresponding to the highest  $q_k(t_k)$  at this step. Now, suppose  $\exists \psi_1, \dots, \psi_\alpha$  such that

$$\mathcal{S}_q(\psi_1, \dots, \psi_i + 1, \dots, \psi_\alpha) = \mathcal{S}_q(\psi_1, \dots, \psi_i, \dots, \psi_\alpha) + q_i(\psi_i + 1)$$

with  $\mathcal{S}_q(\psi_1, \dots, \psi_i + 1, \dots, \psi_\alpha) > \mathcal{S}_q(\phi_1, \dots, \phi_j + 1, \dots, \phi_\alpha)$  and  $n = \psi_1 + \dots + \psi_\alpha = \phi_1 + \dots + \phi_\alpha$ .

By assumption,  $\mathcal{S}_q(\psi_1, \dots, \psi_i, \dots, \psi_\alpha) \leq \mathcal{S}_q(\phi_1, \dots, \phi_j, \dots, \phi_\alpha)$ , it implies that  $q_i(\psi_i + 1) > q_j(\phi_j + 1)$ . But as  $q_j(\phi_j + 1)$  has been selected before  $q_i(\psi_i + 1)$ , so we have  $q_j(\phi_j + 1) > q_i(\psi'_i)$  with  $\psi'_i \leq \psi_i + 1$ , as  $q_i$  is a decreasing function, it comes that  $q_i(\psi_i + 1) \leq q_j(\phi_j + 1)$ , in contradiction with the previous statement.  $\square$

### 5.2.2 Numerical measures

Our practical problem can be related to this lemma considering that  $q(i) = \frac{\#tp \in \mathcal{W}_i}{g} = \frac{\chi_i}{g}$ . But it is important to notice that  $\chi_i$  may not be decreasing functions during a real-life aggregation experiment. We therefore examine how close we are to the hypotheses of the lemma by examining if  $\forall i$ ,  $\chi_i$  are decreasing functions. Note that having decreasing functions  $\chi_i$  means that our classifying features are chosen so that highest ranked pairs have an increased probability of being true positive predictions – which is expected from a *good* classifier.

As shown on the example of Figure 1, the decreasing condition is nearly fulfilled during the process if  $g$  is large enough<sup>5</sup>. However, it does not mean that the best solution for the merging problem is to take  $g$  as large as possible, as the larger  $g$  gets, the lesser  $\chi_i$  is an accurate estimation of the probability that the next pair selected is a true positive prediction. We therefore have to manage a trade-off by tuning parameter  $g$  in order to obtain the best possible performance of *RankMerging*.

---

<sup>5</sup>This example has been extracted from the experiments on the PSP dataset (see later).




Figure 1: Function  $q(i) = \frac{\lambda_i}{g}$  corresponding to ranking  $s_{CN,w}$  for  $g = 100, 1000, 5000$  during a merging process on the PSP dataset.

## 6 Experiments

### 6.1 Benchmarks for comparison

In order to assess the efficiency of *RankMerging*, we compare its performances to existing techniques. First, we compare to the various unsupervised metrics described above. We also compare to classic supervised techniques for classification tasks. We restricted ourselves to several computationally efficient methods, namely nearest neighbors (*NN*), classification trees (*CT*) and *AdaBoost* (*AB*). We have used implementations from Python scikit learn toolkit<sup>6</sup>. These techniques are not specifically designed for ranking tasks, so we obtain different points in the precision-recall space varying the algorithms parameters, respectively the number of neighbors (*NN*), the minimum size of a leaf (*CT*), and the number of trees (*AB*). We recall that usual supervised learning-to-rank techniques are not available in our case.

### 6.2 Series 1

In the first series of experiments, we focus on the PSP dataset and explore the impact of the parameter  $g$  value and the question of the feature selection

<sup>6</sup><http://scikit-learn.org/> , for more details, see [28].



in details. As PSPs only have access to records involving their own clients, they have an incomplete view of the social network as a whole. In several practical applications, however, this missing information is crucial. An important example is churn prediction, that is the detection of clients at risk of leaving to another PSP. This risk is known to depend on the local structure of the social network [10, 26]. Therefore, the knowledge of connections between subscribers of other providers is important for the design of efficient customer retention campaigns. We design this series of experiments for this application, that is to say the prediction of links among the users of a PSP competitor.

### 6.2.1 Protocol

A PSP is usually confronted to the following situation: it has full access to the details of phone calls between its subscribers, as well as between one of its subscriber and a subscriber of another PSP. However, connections between subscribers of other PSPs are hidden. In order to simulate this situation from our dataset, we divide the set of nodes  $\mathcal{V}$  of our data into three sets:  $\mathcal{V}_1$ ,  $\mathcal{V}_2$  and  $\mathcal{V}_3$ , and a partition of the set of links  $\mathcal{E}$  will ensue. During the learning phase, links  $\mathcal{V}_1$ – $\mathcal{V}_1$  and  $\mathcal{V}_1$ – $\mathcal{V}_2$  are known, defining the set of links  $\mathcal{E}_{learn}$  of the graph  $\mathcal{G}_{learn}$ , and we calibrate the model by guessing links  $\mathcal{V}_2$ – $\mathcal{V}_2$ , defining the set  $\mathcal{E}_{cal}$ . During the test phase, all the links are known except for links  $\mathcal{V}_3$ – $\mathcal{V}_3$ , that we guess to evaluate the performances, we denote this set of links  $\mathcal{E}_{perf}$ . Users have been assigned to  $\mathcal{V}_1$ ,  $\mathcal{V}_2$  and  $\mathcal{V}_3$  according to the proportions 50, 25, 25%. With these notations:

- $\mathcal{G}_{learn} = (\mathcal{V}_1 \cup \mathcal{V}_2, \mathcal{E}_{learn})$ , it contains 848,911 nodes and 597,538 links and we aim at predicting the 49,731 links in  $\mathcal{E}_{cal}$ .
- $\mathcal{G}_{test} = (\mathcal{V}, \mathcal{E} \setminus \mathcal{E}_{perf})$ , it contains 1,131,049 nodes and 746,202 links and we aim at predicting the 49,663 links in  $\mathcal{E}_{perf}$ .

Notice that according to our simulation, the learning set and the test set are derived from the same distributions, while it could not be the case in situations involving real world PSPs. However, this assumption is fair to compare the performances of our method to other prediction techniques.

### 6.2.2 Unsupervised learning process

We plot the results obtained on  $\mathcal{G}_{learn}$  to predict  $\mathcal{E}_{cal}$  links for the above classifiers. For the sake of readability, we only represent a selection of them on Figure 2. The evolution of the F-score significantly varies from one classifier to another. For example, it increases sharply for  $CN$ , and then more slowly until reaching its maximum, while  $RWR_w$  rises smoothly before slowly declining. Borda’s aggregation improves the performance of the classification, especially considering the precision on the top-ranked pairs. *RankMerging* method aims at exploiting the differences between these profiles. Given the difficulty of the task, precision is low on average: for instance, when recall is greater than 0.06, precision is lower than 0.3 for all rankers. We used here structural features, making it impossible to predict links between nodes which are too distant from each other. We are therefore limited to small recall values, as increasing the number of predictions lowers dramatically the precision.

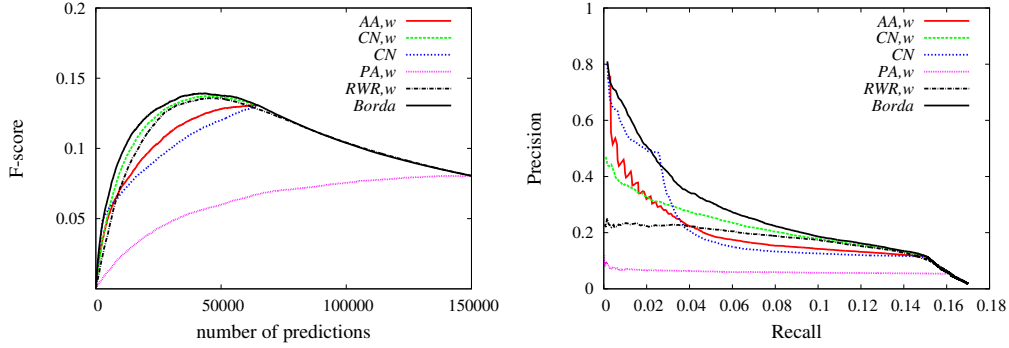


Figure 2: Results obtained on the learning set for various structural classifiers. Left: F-score as a function of the number of predictions. Right: Precision versus Recall curves.

### 6.2.3 Supervised merging process

According to the description in 5.1,  $\phi_i$  are computed on  $\mathcal{G}_{learn}$  to discover  $\mathcal{E}_{cal}$  links, and then used to merge rankings on  $\mathcal{G}_{test}$  to discover  $\mathcal{E}_{perf}$  links, applying the scaling factor  $f \approx 1.5$  to adapt the  $\phi_i$  learnt to the test rankings. Cross-validation is made through a simple hold-out strategy. We will argue

in the following that the user may aggregate as many rankings as possible: the information provided by different rankings may be redundant, but the merging process is such that the addition of a supplementary ranking is computationally cheap, and if a ranking does not bring additional information, it should simply be ignored during the learning process. As of the value of  $g$ , our numerical experiments show that the performance of the algorithm is robust over a sufficiently large range of values (see Table 3), and we extrapolate the best  $g$  value on  $\mathcal{G}_{learn}$  for the aggregation on  $\mathcal{G}_{test}$ .

#### 6.2.4 Results

We plot on Figure 3 the evolution of the F-score and the precision-recall curve obtained with *RankMerging*, for  $g = 200$ , aggregating the rankings of the following classifiers:  $AA_w$ ,  $CN_w$ ,  $CN$ ,  $SR_w$ ,  $Katz_w$  ( $\beta = 0.01$ ),  $PA_w$ ,  $RWR_w$  ( $p = 0.8$ ) and Borda’s method applied to the seven former ones. We observe that *RankMerging* performs better than Borda, and consequently better than all the unsupervised methods explored, especially for intermediary recall values. It was expected, as *RankMerging* incorporates the information of Borda’s aggregation here. We measure the area under the Precision-Recall curves to quantify the performances with a scalar quantity. *RankMerging* increases the area by 8.3% compared to Borda. Concerning the supervised benchmarks, we observe that they perform well, but only for a low number of predictions (comparable to Borda for approximately 1000 to 2000 predictions). Unsurprisingly, *AdaBoost* is an ensemble method and outperforms *Decision Trees* and *Nearest Neighbors* for an optimal parameter choice, but the performances are of the same order of magnitude, in line with the observations in [1]. As formerly stated, these methods are not designed to cover a wide range of the precision-recall space, and therefore perform very poorly out of their optimal region of use.

On the minus side, *RankMerging* has been designed for classification problems with large number of predictions. The window size  $g$  implies an averaging effect which causes the method to lack efficiency on the top-ranked items, as can be seen on Fig. 3. As a consequence, it is not suited to problems with low number of predictions, as it is often the case for information retrieval tasks for example.

**Feature selection.** We evaluate the influence of the structural metrics in Table 3. A comprehensive study of the matter would be long and repetitive

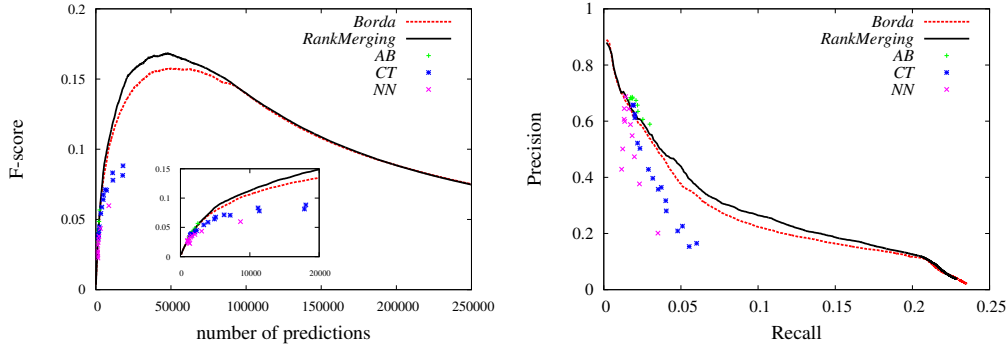


Figure 3: Results obtained with *RankMerging* on the test set ( $g = 200$ ), compared to various benchmarks. Left: F-score as a function of the number of predictions. Right: Precision versus Recall curves.

considering all possible combinations of rankings, and we restrict ourselves to a few examples. An important result to notice here is that the addition of a ranking does not decrease the quality of the merging process – except for small fluctuations. So ideally, a user may add any ranking to his own implementation of the method, whatever the source of information is, it should improve – or at least not worsen – the prediction. Note that the design of the method is supposed to favour this behaviour: if a ranking brings too little information to the aggregation process, then it should be ignored during the learning process. The problem of feature selection is critical for the link prediction problem, for instance consensus methods performances drop dramatically with a poor feature choice, so these experimental observations are encouraging for future uses of *RankMerging*.

**Parameter  $g$ .** The dependency on  $g$  value is shown on Table 4. Results indicate that the performances are close to the maximum within the interval  $[100; 300]$  on both the learning and test sets. This observation suggests the possibility of tuning  $g$  in the testing phase from the values of  $g$  during the learning process. It is interesting to note that the performance is maximum for these intermediate values, around  $g = 200$ . This is the right balance between small  $g$  which fluctuate too much to provide a good local optimum, and large  $g$ , which average the quality of rankings on a too large window.

$AA_w$	$CN_w$	$CN$	$SR_w$	$Katz_w$	$PA_w$	$RWR_w$	$Borda$	imp.(%)
X	X	X	X	X	X	X	X	8.3
X	X	X	X	X	X	X		3.2
X	X	X	X	X	X			-0.7
X	X	X	X	X				-1.0
X	X	X	X					-2.0
	X	X	X	X	X	X	X	8.2
		X	X	X	X	X	X	8.1
			X	X	X	X	X	3.7
				X	X	X	X	3.8

Table 3: Improvement (in %) to Borda’s method of the area under the curve in the precision-recall space, for the aggregation of different rankings.

$g$	10	100	200	300	400	500	1000	2000
imp.	-0.8	5.5	5.4	5.2	5.0	4.7	4.0	2.7

$g$	10	100	200	300	400	500	1000	2000
imp.	2.7	8.2	8.3	7.9	7.4	7.2	6.4	5.6

Table 4: Improvement to Borda’s method of the area under the curve in the precision-recall space, for different values of  $g$ . Top: learning set, bottom: test set.

## 6.3 Series 2

In the second series of experiments, we focus on DBLP and Pokec datasets. Here, we investigate in details the impact of the sizes of the learning and test sets.

### 6.3.1 Protocol

A few points differ from the protocol of first series. Here, all the nodes belong to both  $\mathcal{G}_{learn}$  and  $\mathcal{G}_{test}$ , the partition is made on the set of links  $\mathcal{E}$ :

- $\mathcal{E}_{learn}$  are the links of  $\mathcal{G}_{learn}$ ,
- $\mathcal{E}_{cal}$  are the links to guess during the learning phase, to calibrate our model,

- $\mathcal{E}_{perf}$  are the links to guess during the testing phase, to evaluate the performance of the method.

During the learning phase,  $\mathcal{E}_{learn}$  links are used to guess  $\mathcal{E}_{cal}$  links. During the test phase, the links of  $\mathcal{G}_{test}$ , that is to say  $\mathcal{E}_{test} = \mathcal{E}_{learn} \cup \mathcal{E}_{cal}$  are used to guess  $\mathcal{E}_{perf}$ . We generate several samples such that  $|\mathcal{E}_{cal}| = |\mathcal{E}_{perf}|$ , but with various values for  $|\mathcal{E}_{cal}|$ , the missing information increases as this set grows larger. The samples are defined by the missing links ratio  $|\mathcal{E}_{cal}|/|\mathcal{E}|$ .

Another noteworthy difference is that these networks have a much higher average degree than the PSP network, making the computation of the global rankers expensive in both memory and time. We therefore limited ourselves to the less costly local and intermediary metrics, more precisely:  $AA_w$ ,  $CN_w$ ,  $SR_w$ ,  $RA_w$ ,  $LP_w$  ( $\gamma = 0.1$ ) and Borda’s aggregation.

For the same reasons, the number of pairs is very large when only considering nodes at distance 2 (typically larger than  $10^8$ ). Since it would not make much sense to predict so many links in a social network, we choose to reduce the rankings length by focusing on the intersection of rankings: we kept only the top  $10^6$  pairs in the intersection of the different rankings. Borda’s method is then applied on the intersected rankings. In this setting of experiments, the scaling factor  $f = 1$  during the whole series.

### 6.3.2 Results

We plot on Figure 4 the results obtained on the test sets of both DBLP and Pokec datasets, for samples where the missing links ratio  $|\mathcal{E}_{cal}|/|\mathcal{E}| = 0.077$ . In both cases, *RankMerging* outperforms the unsupervised methods, but the improvement is much more visible for DBLP than for Pokec.

In the case of DBLP, a closer examination shows that at the beginning of the process, we closely follow Sørensen index curve, which is the most efficient unsupervised ranking for low number of predictions. Then its performance drops, and other rankings are chosen during the aggregation process (mostly Local Path ranking). In the case of Pokec, pairs aggregated initially mainly come from Adamic-Adar ranking, then this index soon gets less efficient than Resource Allocation index, which takes over until the end of the merging process. Pokec gives a good indication of what happens when the rankings are not complementary enough, and that one of them is more efficient than the others: the aggregation nearly always choose the pairs from this ranking and do not improve significantly the performance. Notice that in both cases,

Borda’s method is not the most efficient unsupervised method: as some rankings perform very poorly, they dramatically hinder the performance of Borda’s method.

The comparison to supervised classification methods is interesting here<sup>7</sup>: *Classification Trees* as well as *Nearest Neighbors* perform poorly on Pokec, however the results on DBLP show that *NN* and above all *CT* perform better than *RankMerging* on a limited range. This observation highlights a limit of our learning-to-rank aggregation method for the purpose of link prediction: the prediction performance cannot be largely better than the performances of the rankings that have been used as features for the learning, while classification methods can, the compensation being the severe constraints on the number of predictions.

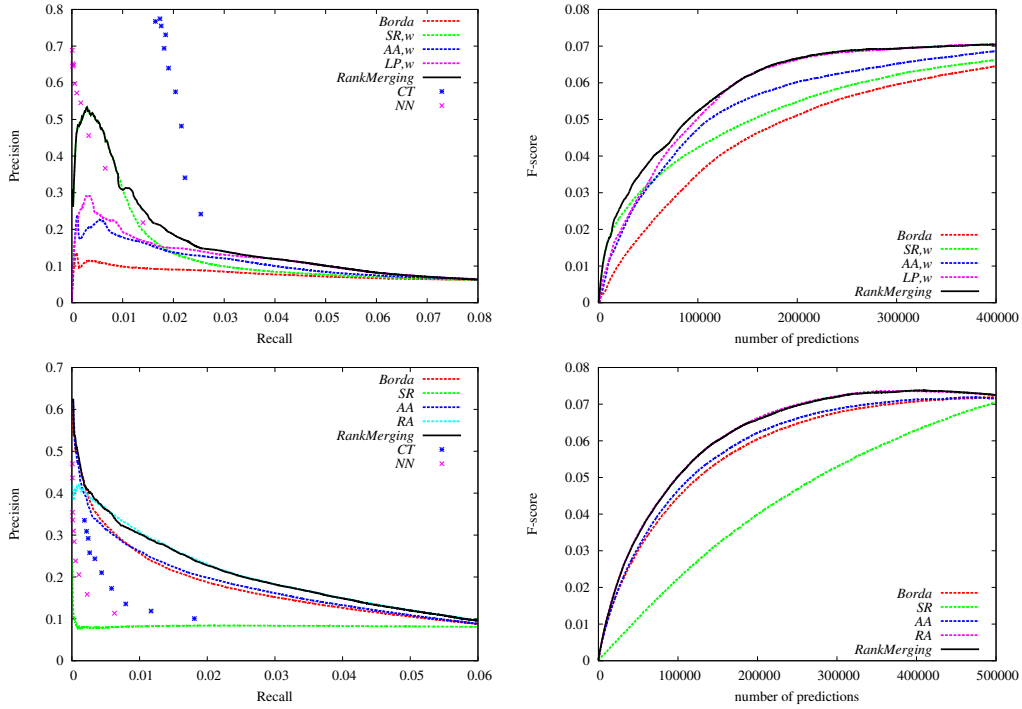


Figure 4: Results obtained with *RankMerging* on the test set. Top: DBLP ( $g = 4000$ ). Bottom: Pokec ( $g = 2000$ ).

<sup>7</sup>Notice that *AdaBoost* does not provide exploitable results on both networks as it predicts very few links.

**Learning set size.** We now explore the impact of the learning set size on the performance of the method. We generated five different samples with missing links ratios: 0.05, 0.10, 0.15, 0.20 and 0.25. On Figure 5, we plot the area under the Precision-Recall curves for the samples generated, and compare them to the most efficient unsupervised methods tested. We observe that *RankMerging* outperforms the unsupervised methods in nearly all cases, but the differences vanish when the size of the learning set decrease, that is to say when the missing information grows. It seems that this observation stems from the fact that a ranker dominate the others when information is missing, so that the merging method tend to stick to the most performing metric.




Figure 5: Area under the precision-recall curve for different missing links ratios. Left: DBLP dataset, right: Pokec dataset.

**Experimental running times.** To give the reader a better grasp of the practical running times, we indicate briefly the order of magnitudes of the different experiments. Of course, these running times depend on the implementations chosen. We used standard ones on the same machine throughout the experiments (a workstation with  $16 \cdot 3$  GHz CPU). Unsupervised rankings production step costs from 10 to 1000 minutes (depending on the dataset and the scoring index), while the whole merging process is shorter than a minute for all the datasets.



## 7 Conclusion

In this work, we have presented *RankMerging*, a supervised machine learning framework which combines rankings from unsupervised classifiers to improve the performance of link prediction. This learning-to-rank method is straightforward and computationally cheap – its complexity is  $O(n.\alpha)$ , where  $\alpha$  is the number of rankings aggregated and  $n$  the number of predictions. It is suited to prediction in social networks, as  $n$  can be tuned according to the application’s needs. As we have discussed the values of  $n$  are generally large, that is our method is designed to yield a large number of predictions. Indeed, the precision on top-ranked items is not significantly improved, making the method ineffective for most information retrieval purposes.

So far, we have exclusively focused on structural information in order to predict unknown links. However, the framework is general and any feature providing a ranking for likely pairs of nodes can be incorporated. Additional structural classifiers are an option, but other types of attributes can also be considered, such as the profile of the users (age, hometown, etc), or timings of their interactions. In the latter case, for instance, if  $i$  and  $j$  are both interacting with  $k$  within a short span of time, it is probably an indication of a connection between  $i$  and  $j$ . From a theoretical perspective, *RankMerging* provides a way to uncover the mechanisms of link creation, by identifying which sources of information play a dominant role in the quality of a prediction. The method could be applied to other types of networks, especially when links are difficult to detect. Applications include network security – for example by detecting the existence of connections between machines of a botnet – and biomedical engineering – for screening combinations of active compounds and experimental environments in the purpose of medicine discovery.

## Acknowledgements

The authors would like to thank Emmanuel Viennet and Maximilien Danisch for useful bibliographic indications. This paper presents research results of the Belgian Network DYSCO (Dynamical Systems, Control, and Optimization), funded by the Interuniversity Attraction Poles Programme, initiated by the Belgian State, Science Policy Office. The scientific responsibility rests with its authors. We also acknowledge support from FNRS and the European

## References

- [1] M. Al Hasan, V. Chaoji, S. Salem, and M. Zaki. Link prediction using supervised learning. In SDM'06: Workshop on Link Analysis, Counter-terrorism and Security, 2006.
- [2] L. Backstrom and J. Leskovec. Supervised random walks: predicting and recommending links in social networks. In Proceedings of the fourth ACM international conference on Web search and data mining, pages 635–644. ACM, 2011.
- [3] N. Benchettara, R. Kanawati, and C. Rouveirol. Supervised machine learning applied to link prediction in bipartite social networks. In International Conference on Advances in Social Networks Analysis and Mining (ASONAM), pages 326–330. IEEE, 2010.
- [4] C.A. Bliss, M.R. Frank, C.M. Danforth, and P.S. Dodds. An evolutionary algorithm approach to link prediction in dynamic social networks. arXiv:1304.6257, 2013.
- [5] C.J.C. Burges, K.M. Svore, P.N. Bennett, A. Pastusiak, and Q. Wu. Learning to rank using an ensemble of lambda-gradient models. Journal of Machine Learning Research-Proceedings Track, 14:25–35, 2011.
- [6] Z. Cao, T. Qin, T.Y. Liu, M.F. Tsai, and H. Li. Learning to rank: from pairwise approach to listwise approach. In Proceedings of the 24th international conference on Machine learning, pages 129–136. ACM, 2007.
- [7] O. Chapelle, Y. Chang, and T.Y. Liu. Future directions in learning to rank. In Yahoo! Learning to Rank Challenge, pages 91–100, 2011.
- [8] O. Chapelle and S.S. Keerthi. Efficient algorithms for ranking with svms. Information Retrieval, 13(3):201–215, 2010.
- [9] P.M. Comar, P.N. Tan, and A.K. Jain. Linkboost: A novel cost-sensitive boosting framework for community-level network link prediction. In

- 11th International Conference on Data Mining (ICDM), pages 131–140. IEEE, 2011.
- [10] K. Dasgupta, R. Singh, B. Viswanathan, D. Chakraborty, S. Mukherjea, A.A. Nanavati, and A. Joshi. Social ties and their relevance to churn in mobile telecom networks. In Proceedings of the 11th International Conference on Extending Database Technology, pages 668–677. ACM, 2008.
  - [11] D. Davis, R. Lichtenwalter, and N.V. Chawla. Supervised methods for multi-relational link prediction. Social Network Analysis and Mining, 3(2):127–141, 2013.
  - [12] J.C. de Borda. Mémoire sur les élections au scrutin. 1781.
  - [13] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In Proceedings of the 10th international conference on World Wide Web, pages 613–622. ACM, 2001.
  - [14] Y. Freund, R. Iyer, R.E. Schapire, and Y. Singer. An efficient boosting algorithm for combining preferences. The Journal of machine learning research, 4:933–969, 2003.
  - [15] R. Herbrich, T. Graepel, and K. Obermayer. Large margin rank boundaries for ordinal regression. Advances in neural information processing systems, pages 115–132, 1999.
  - [16] Z. Huang, X. Li, and H. Chen. Link prediction approach to collaborative filtering. In Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries, pages 141–142. ACM, 2005.
  - [17] H. Kashima, T. Kato, Y. Yamanishi, M. Sugiyama, and K. Tsuda. Link propagation: A fast semi-supervised learning algorithm for link prediction. In SDM, volume 9, pages 1099–1110. SIAM, 2009.
  - [18] G. Kossinets and D.J. Watts. Empirical analysis of an evolving social network. Science, 311(5757):88–90, 2006.
  - [19] J. Leskovec, L. Backstrom, R. Kumar, and A. Tomkins. Microscopic evolution of social networks. In Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining, pages 462–470. ACM, 2008.

- [20] D. Liben-Nowell and J. Kleinberg. The link-prediction problem for social networks. Journal of the American society for information science and technology, 58(7):1019–1031, 2007.
- [21] R.N. Lichtenwalter, J.T. Lussier, and N.V. Chawla. New perspectives and methods in link prediction. In Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining, pages 243–252. ACM, 2010.
- [22] T.Y. Liu. Learning to rank for information retrieval. Foundations and Trends in Information Retrieval, 3(3):225–331, 2009.
- [23] L. Lü and T. Zhou. Link prediction in complex networks: A survey. Physica A: Statistical Mechanics and its Applications, 390(6):1150–1170, 2011.
- [24] A.K. Menon and C. Elkan. Link prediction via matrix factorization. In Machine Learning and Knowledge Discovery in Databases, pages 437–452. Springer, 2011.
- [25] T. Murata and S. Moriyasu. Link prediction of social networks based on weighted proximity measures. In International Conference on Web Intelligence, pages 85–88. IEEE, 2007.
- [26] B. Ngonmang, E. Viennet, and M. Tchuente. Churn prediction in a real online social network using local community analysis. In Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012), pages 282–288. IEEE Computer Society, 2012.
- [27] M. Pavlov and R. Ichise. Finding experts by link prediction in co-authorship networks. FEWS, 290:42–55, 2007.
- [28] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. Journal of Machine Learning Research, 12:2825–2830, 2011.
- [29] M. Pujari and R. Kanawati. Supervised rank aggregation approach for link prediction in complex networks. In Proceedings of the 21st

- international conference companion on World Wide Web, pages 1189–1196. ACM, 2012.
- [30] T. Raeder, O. Lizardo, D. Hachen, and N.V. Chawla. Predictors of short-term decay of cell phone contacts in a large scale communication network. Social Networks, 33(4):245–257, 2011.
  - [31] S. Scellato, A. Noulas, and C. Mascolo. Exploiting place features in link prediction on location-based social networks. In Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining, pages 1046–1054. ACM, 2011.
  - [32] D. Sculley. Rank aggregation for similar items. In SDM, pages 587–592. SIAM, 2007.
  - [33] T. Tyenda, R. Angelova, and S. Bedathur. Towards time-aware link prediction in evolving social networks. In Proceedings of the 3rd Workshop on Social Network Mining and Analysis, page 9. ACM, 2009.
  - [34] T. Zhou, L. Lü, and Y.C. Zhang. Predicting missing links via local information. The European Physical Journal B, 71(4):623–630, 2009.